

The configuration space of a robotic arm in a tunnel.

Federico Ardila*

Hanner Bastidas†

Cesar Ceballos‡

John Guo§

Abstract

We study the motion of a robotic arm inside a rectangular tunnel. We prove that the configuration space of all possible positions of the robot is a CAT(0) cubical complex. This allows us to use techniques from geometric group theory to find the optimal way of moving the arm from one position to another. We also compute the diameter of the configuration space, that is, the longest distance between two positions of the robot.

1 Introduction

We consider a robotic arm $R_{m,n}$ of length n moving in a rectangular tunnel of width m without self-intersecting. The robot consists of n links of unit length, attached sequentially, and facing up, down, or right. The base is affixed to the lower left corner. Figure 1 illustrates a possible position of the arm of length 8 in a tunnel of width 2.



Figure 1: The robotic arm $R_{8,2}$.

The robotic arm may move freely using two kinds of local moves:

- *Switch corners*: Two consecutive links facing different directions interchange directions.
- *Flip the end*: The last link of the robot rotates 90° without intersecting itself.




Figure 2: The two kinds of local moves of the robotic arm.

We study the following fundamental problem.

Problem 1.1. *Move the robotic arm $R_{m,n}$ from one given position to another optimally.*

*San Francisco State University, San Francisco, USA; Univ. de Los Andes, Bogotá, Colombia; federico@sfsu.edu.

†Departamento de Matemáticas, Universidad del Valle, Cali, Colombia; hanner.bastidas@correounivalle.edu.co.

‡Faculty of Mathematics, University of Vienna, Austria; cesar.ceballos@univie.ac.at.

§San Francisco State University, San Francisco, USA; jguo@mail.sfsu.edu.

This project is part of the SFSU-Colombia Combinatorics Initiative. FA was supported by the US NSF CAREER Award DMS-0956178 and NSF Combinatorics Award DMS-1600609. CC was supported by a Banting Postdoctoral Fellowship of the Canadian government, a York University research grant, and the Austrian FWF Grant F 5008-N15.

When we are in a city we do not know well and we are trying to get from one location to another, we will usually consult a map of the city to plan our route. This is a simple but powerful idea. Our strategy to approach Problem 1.1 is pervasive in many branches of mathematics: we will be to build and understand the “map” of all possible positions of the robot; we call it the *configuration space* $\mathcal{S}_{m,n}$. Such spaces are also called *state complexes*, *moduli spaces*, or *parameter spaces* in other fields. Following work of Abrams–Ghrist [1] in applied topology and Reeves [11] in geometric group theory, Ardila, Owen, and Sullivan [5] and Ardila, Baker, and Yatchak [3] showed how to solve Problem 1.1 for robots whose configuration space is $CAT(0)$; this is a notion of non-positive curvature defined in Section 4. This is the motivation for our main result.

Theorem 1.2. *The configuration space $\mathcal{S}_{m,n}$ of the robotic arm $R_{m,n}$ of length n in a tunnel of width w is a $CAT(0)$ cubical complex.*




Figure 3: The configuration space $\mathcal{S}_{6,2}$ of the robotic arm $R_{6,2}$.

In light of [3], Theorem 1.2 provides a solution to Problem 1.1.

Corollary 1.3. *There is an explicit algorithm, implemented in Python, to move the robotic arm $R_{m,n}$ optimally from one given position to another.*

The reader may visit <http://math.sfsu.edu/federico/Articles/movingrobots.html> to watch a video preview of this program or to download the Python code.

Abrams and Ghrist showed that cubical complexes appear in a wide variety of contexts where a discrete system moves according to local rules – such as domino tilings, reduced words in a Coxeter group, and non-intersecting particles moving in a graph – and that these cubical complexes are sometimes $CAT(0)$ [1]. The goal of this paper is to illustrate how the techniques of [3, 5] may be used to prove that configuration spaces are $CAT(0)$, uncover their underlying combinatorial structure, and navigate them optimally.

To our knowledge, Theorem 1.2 provides one of the first infinite families of configuration spaces which are $CAT(0)$ cubical complexes, and the first one to invoke the full machinery of *posets with inconsistent pairs* of [3, 5]. It is our hope that the techniques employed here will provide a blueprint for the study of $CAT(0)$ cubical complexes in many other contexts.

We now describe the structure of the paper. In Section 2 we define more precisely the configuration space $\mathcal{S}_{m,n}$ of the robotic arm $R_{m,n}$. Section 3 is devoted to collecting some preliminary enumerative evidence for our main result, Theorem 1.2. It follows from very general results of Abrams and Ghrist [1] that the configuration space $\mathcal{S}_{m,n}$ is a cubical complex. Also, we know from work of Gromov [9] that $\mathcal{S}_{m,n}$ will be $CAT(0)$ if and only if it is contractible.

Before launching into the proof that $\mathcal{S}_{m,n}$ is CAT(0), we first verify in the special case $m = 2$ that this space has the correct Euler characteristic, equal to 1. We do it as follows.

Theorem 1.4. *If $c_{n,d}$ denotes the number of d -dimensional cubes in the configuration space $\mathcal{S}_{2,n}$ for a robot of width 2, then*

$$\sum_{n,d \geq 0} c_{n,d} x^n y^d = \frac{1 + x^2 + 2x^3 - x^4 + xy + x^2y + 4x^3y + x^4y + x^3y^2 + 2x^4y^2 + x^5y^2}{1 - 2x + x^2 - x^3 - x^4 - 2x^4y - 2x^5y - x^5y^2 - x^6y^2}.$$

The Euler characteristic of $\mathcal{S}_{2,n}$ is $\chi(\mathcal{S}_{2,n}) = c_{n,0} - c_{n,1} + \dots$. Substituting $y = -1$ above we find, in an expected but satisfying miracle of cancellation, that the generating function for $\chi(\mathcal{S}_{2,n})$ is $1/(1 - x) = 1 + x + x^2 + \dots$. We obtain:

Theorem 1.5. *The Euler characteristic of the configuration space $\mathcal{S}_{2,n}$ equals 1.*

In Section 4 we collect the tools we use to prove Theorem 1.2. Ardila, Owen, and Sullivan [5] gave a bijection between CAT(0) cubical complexes X and combinatorial objects $P(X)$ called *posets with inconsistent pairs* or PIPs. The PIP $P(X)$ is usually much simpler, and serves as a “remote control” to navigate the space X . This bijection allows us to prove that a (rooted) cubical complex is CAT(0) by identifying its corresponding PIP.¹

We use this technique to prove Theorem 1.2 in Section 5. To do this, we introduce the *coral PIP* $C_{m,n}$ of *coral tableaux*, illustrated in Figure 4. By studying the combinatorics of coral tableaux, we show the following:

Proposition 1.6. *The coral PIP $C_{m,n}$ is the PIP that certifies that the configuration space $\mathcal{S}_{m,n}$ is a CAT(0) cubical complex.*




Figure 4: The coral PIPs (“remote controls”) for the robotic arms of length 1, 2, 3, 4, 5, 6, 9 in a tunnel of width 2.

In Section 6 we use the combinatorics of coral tableaux to find a combinatorial formula for the distance between any two positions of the robotic arm. This allows us to find the longest of such distances – that is, the diameter of the configuration space $\mathcal{S}_{m,n}$ in the l_1 -metric – answering a question of the first author. [2]

¹It is worth noting that PIPs are known as *event systems* [15] in computer science and are closely related to *pocsets* [12, 13] in geometric group theory.

Theorem 1.7. *The diameter of the transition graph of the robot $R_{m,n}$ of length n in a tunnel of width $m \geq 2$ is*

$$\text{diam } G(R_{m,n}) = \begin{cases} d(L_{m,n}, H_{m,n}) & \text{for } n < 6 \\ d(L_{m,n}, L_{m,n}^+) & \text{for } n \geq 6, \end{cases}$$

where $L_{m,n} = u^m r d^m r u^m r \dots$ is the left justified position, $L_{m,n}^+ = (urdr)L_{m,n-4}$ and $H_{m,n}$ is the fully horizontal position (see Figure 19). A precise formula is stated in Theorem 6.4.

Finally, in Section 7 we discuss our solution to Problem 1.1. As explained in [3], we use the coral PIP as a remote control for our robotic arm. This allows us to implement an algorithm to move the robotic arm $R_{m,n}$ optimally from one position to another.

2 The configuration space

Recall that $R_{m,n}$ is a robotic arm consisting of n links of unit length attached sequentially, and facing up, down, or right. The robot is inside a rectangular tunnel of width m and pinned down at the bottom left corner of the tunnel. It moves inside the tunnel without self-intersecting, by switching corners and flipping the end as illustrated in Figure 2.

In this section we describe the configuration space $\mathcal{S}_{m,n}$ of all possible positions of the robot $R_{m,n}$. We begin by considering the *transition graph* $G(R_{m,n})$ whose vertices are the possible states of the robot, and whose edges correspond to the allowable moves between them. Figure 5 illustrates the transition graph $G(R_{2,4})$ of a robotic arm of length 4.




Figure 5: The transition graph of a robotic arm of length 4.

As these examples illustrate, each one of these graphs is the 1-skeleton of a cubical complex. For example, consider a position u which has two legal moves a and b occurring in disjoint parts of the arm. We call a and b *physically independent* or *commutative* because $a(b(u)) = b(a(u))$. In this case, there is a square connecting $u, a(u), b(a(u)) = a(b(u))$, and $b(u)$ in $G(R_{m,n})$.

More generally, we obtain the configuration space by filling all the cubes of various dimensions that we see in the transition graph. Let us make this precise.

Definition 2.1. *The configuration space or state complex $\mathcal{S}_{m,n}$ of the robot $R_{m,n}$ is the following cubical complex. The vertices correspond to the states of $R_{m,n}$. An edge between vertices u and v corresponds to a legal move which takes the robot between positions u and v . The k -cubes correspond to k -tuples of commutative moves: Given k such moves which are applicable at a state u , we can obtain 2^k different states from u by performing a subset of these k moves; these are the vertices of a k -cube in $\mathcal{S}_{m,n}$. We endow $\mathcal{S}_{m,n}$ with a Euclidean metric by letting each k -cube be a unit cube.*

Figure 3 shows the configuration space $\mathcal{S}_{2,6}$; every square and cube in the diagram is filled in.

3 Face enumeration and the Euler characteristic of $\mathcal{S}_{2,n}$

The main structural result of this paper, Theorem 1.2, is that the configuration space $\mathcal{S}_{m,n}$ of our robot is a *CAT(0) cubical complex*. This is a subtle metric property defined in Section 4 and proved in Section 5. As a prelude, this section is devoted to proving a partial result in that direction. It is known [7, 9] that CAT(0) spaces are contractible, and hence have Euler characteristic equal to 1. We now prove:

Theorem 1.5. The Euler characteristic of the configuration space $\mathcal{S}_{2,n}$ equals 1.

This provides enumerative evidence for our main result in width $m = 2$. While we were stuck for several weeks trying to prove Theorem 1.2, we found this evidence very encouraging.

To prove Theorem 1.5, our strategy is to compute the f -vector of $\mathcal{S}_{2,n}$. Recall that the f -vector of a d -dimensional polyhedral complex X is $f_X = (f_0, f_1, \dots, f_d)$ where f_k is the number of k -dimensional faces. The Euler characteristic of X is $\chi(X) = f_0 - f_1 + \dots + (-1)^d f_d$. Table 1 shows the f -vectors of the cubical complexes of the robotic arms of length $n \leq 6$. For instance, the complex of Figure 3 contains 53 vertices, 81 edges, 30 squares, and 1 cube. We now carry out this computation for all n .

n	f_0	f_1	f_2	f_3	$\chi(\mathcal{S}_{2,n})$
1	2	1	0	0	1
2	4	3	0	0	1
3	8	8	1	0	1
4	15	18	4	0	1
5	28	38	11	0	1
6	53	81	30	1	1

Table 1: The f -vectors of the cubical complexes $\mathcal{S}_{2,n}$ for arms of length $n \leq 6$.

3.1 Face enumeration

We now compute the generating function for the f -vectors of the configuration spaces $\mathcal{S}_{2,n}$. We proceed in several steps.

3.1.1 Cubes and partial states

Consider a d -cube in the configuration space $\mathcal{S}_{2,n}$; it has 2^d vertices. If we superimpose the corresponding 2^d positions of the robotic arm, we obtain a sequence of edges, squares, and possibly a “claw” in the last position, as illustrated in Figure 6. The number of squares (including the claw if it is present) is d , corresponding to the d physically independent moves being represented by this cube. We call the resulting diagram a *partial state*, and let its *weight* be $x^n y^d$. The partial states of weight $x^n y^d$ are in bijection with the d -cubes of $\mathcal{S}_{2,n}$.




Figure 6: A partial state corresponding to a 6-cube in the configuration space $\mathcal{S}_{2,20}$.

Each partial state gives rise to a word in the alphabet $\{r, v, \ell, \square, \perp\}$, where:

- r represents a horizontal link of the robot facing to the right. Its weight is x .
- v represents a vertical link. Its weight is x .
- \square represents a square, which comes from a move that switches corners of two consecutive links facing different directions. Its weight is x^2y .



- \perp represents a claw, which comes from a move that flips the end of the robot, with the horizontal link facing to the right. Its weight is xy .



For example, the partial state of Figure 6 gives rise to the word $w = r\square\square rv\square r\square rr\square rr v\perp$. The weight of the partial state is the product of the weights of the individual symbols; in this case it is $x(x^2y)(x^2y)xx(x^2y)x(x^2y)xxx(xy) = x^{20}y^6$. It is worth remarking that this word does not determine the partial state uniquely; the reader is invited to construct another state giving rise to the same word w above.

3.1.2 Factorization of partial states into irreducibles

Our next goal is to use generating functions to enumerate all *partial states* according to their length and dimension. The key idea is that we can “factor” a partial state uniquely as a concatenation of irreducible factors. Each time the partial state enters one of the borders of the tunnel, we start a new factor. For example, the factorization of the partial state of Figure 6 is shown in Figure 7.




Figure 7: The partial state of Figure 6 has a factorization of the form $M_1M_5M_1F_7$. (See Tables 2 and 3.)

Definition 3.1. Let P be the set of all partial states of robotic arms in a tunnel of width 2.

- A partial state of the robot is called *irreducible* if
 - its first step is a horizontal link along the bottom border of the tunnel, and
 - its final step is vertical or square, and is its first arrival at the same or opposite border.
- A partial state of the robot is called *irreducible final* if it is empty or
 - its first step is a horizontal link along the bottom border of the tunnel, and
 - either its final step is a claw which is also its first arrival at the same or opposite border, or it never arrives at the same or the opposite border.

Let M and F be the sets of irreducible and irreducible final partial states.

Let $\mathcal{S} = \bigcup_{n=0}^{\infty} \mathcal{S}_{2,n}$ be the disjoint union of the configuration spaces of all robotic arms of all lengths in width 2. Let B^* be the collection of all words that can be made with alphabet B . For instance, $a^* = \{\emptyset, a, aa, aaa, aaaa, \dots\}$ and $\{a, b\}^* = \{\emptyset, a, b, aa, ab, ba, bb, aaa, aab, \dots\}$.

Proposition 3.2. *The partial states in \mathcal{S} starting with a right step r are in weight-preserving bijection with the words in M^*F ; that is, each partial state in \mathcal{S} corresponds to a unique word of the form $m_1m_2\dots m_\ell f$ with $m_i \in M$ and $f \in F$.*

Proof. It is clear from the definitions that every partial state that starts with a horizontal step r factors uniquely as a concatenation $m_1^\pm m_2^\pm \dots m_\ell^\pm f^\pm$ where each $m_i \in M$, $f \in F$, and p^\pm equals p or its reflection p^- across the horizontal axis. It remains to observe that whether m_i^\pm is m_i or m_i^- , and whether p^\pm is either p or p^- , is determined completely by the previous terms of the sequence. \square

Corollary 3.3. *If the generating functions for partial states, irreducible partial states, and irreducible final partial states are $C(x, y)$, $M(x, y)$, $F(x, y)$ respectively, then*

$$1 + xC(x, y) = \frac{F(x, y)}{1 - M(x, y)}.$$

Proof. This follows from Proposition 3.2. The extra factor of x comes from the fact that Proposition 3.2 is counting partial states with an initial right step. \square

3.1.3 Enumeration of irreducible partial states

Let us compute the generating function $M(x, y)$ for irreducible partial states.

Proposition 3.4. *The generating function for the irreducible partial states M is*

$$M(x, y) = \frac{x^3 + x^4 + 2x^4y + 2x^5y + x^5y^2 + x^6y^2}{(1 - x)^2}.$$

Proof. The word of an irreducible partial state has exactly two symbols that contribute a vertical move, which can be either a v or a \square . Thus there are eight different families M_1, \dots, M_8 , corresponding to the irreducible partial states of the following forms:

$$\begin{array}{cccc} \dots \square \dots \square & \dots \square \dots \square' & \dots \square \dots v & \dots \square \dots v' \\ \dots v \dots \square & \dots v \dots \square' & \dots v \dots v & \dots v \dots v' \end{array}$$

where \square' and v' represent a move whose vertical step is in the opposite direction to the previous vertical step. Table 2 illustrates these 8 families together with their corresponding generating functions.

Consider for example the family M_2 consisting of partial states of the form $\dots \square \dots \square'$. We must have at least one horizontal step before the first \square , and at least one horizontal step between the two \square s, to make sure they do not intersect. Therefore the partial states in M_2 are the words in the language $(rr^*)\square(rr^*)\square'$, whose generating function is

$$\left(x \cdot \frac{1}{1 - x}\right) x^2 y \left(x \cdot \frac{1}{1 - x}\right) x^2 y = \frac{x^6 y^2}{(1 - x)^2}.$$

The other formulas in Table 2 follow similarly. Finally, $M(x, y)$ is obtained by adding the eight generating functions in the table. \square

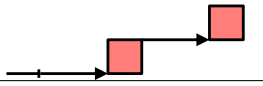
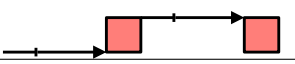
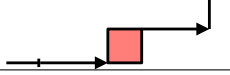
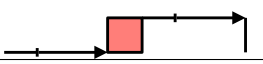
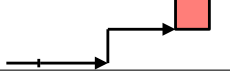
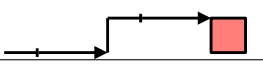
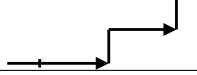
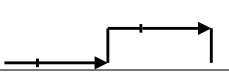
Type	Illustration	Generating function
$M_1 = (rr^*)\square(r^*)\square$		$\frac{x^5 y^2}{(1-x)^2}$
$M_2 = (rr^*)\square(rr^*)\square'$		$\frac{x^6 y^2}{(1-x)^2}$
$M_3 = (rr^*)\square(r^*)v$		$\frac{x^4 y}{(1-x)^2}$
$M_4 = (rr^*)\square(rr^*)v'$		$\frac{x^5 y}{(1-x)^2}$
$M_5 = (rr^*)v(r^*)\square$		$\frac{x^4 y}{(1-x)^2}$
$M_6 = (rr^*)v(rr^*)\square'$		$\frac{x^5 y}{(1-x)^2}$
$M_7 = (rr^*)v(r^*)v$		$\frac{x^3}{(1-x)^2}$
$M_8 = (rr^*)v(rr^*)v'$		$\frac{x^4}{(1-x)^2}$

Table 2: The eight types of irreducible partial states and their generating functions.

3.1.4 Enumeration of irreducible final partial states

Now let us compute the generating function $F(x, y)$ for irreducible final partial states.

Proposition 3.5. *The generating function for the final irreducible partial states is*

$$F(x, y) = \frac{1 - x + x^2 + x^2 y + x^3 y + x^4 y + x^4 y^2 + x^5 y^2}{(1 - x)^2}.$$

Proof. The word of each irreducible final partial state has at most one symbol among $\{v, \square\}$, and can possibly end with \sqcup . Again, we let \sqcup' represent a move whose vertical step is in the opposite direction to the previous vertical step. Thus there are eight different families F_1, \dots, F_8 , corresponding to the irreducible partial states of the following forms:

$$\begin{array}{ccc} \dots & \dots & \dots \sqcup \\ \dots v \dots & \dots v \dots \sqcup & \dots v \dots \sqcup' \\ \dots \square \dots & \dots \square \dots \sqcup & \dots \square \dots \sqcup' \end{array}$$

Table 3 shows the eight different families of possibilities together with their corresponding generating functions. To obtain $F(x, y)$ we add their eight generating functions. \square

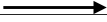
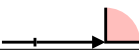
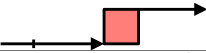
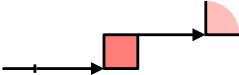
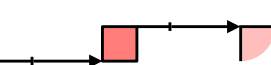


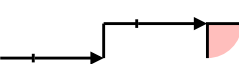
Irreducible move	Illustration	Generating function
$F_1 = r^*$		$\frac{1}{1-x}$
$F_2 = (rr^*)_{\perp}$		$\frac{x^2y}{1-x}$
$F_3 = (rr^*)\square(r^*)$		$\frac{x^3y}{(1-x)^2}$
$F_4 = (rr^*)\square(r^*)_{\perp}$		$\frac{x^4y^2}{(1-x)^2}$
$F_5 = (rr^*)\square(rr^*)_{\perp}'$		$\frac{x^5y^2}{(1-x)^2}$
$F_6 = (rr^*)v(r^*)$		$\frac{x^2}{(1-x)^2}$
$F_7 = (rr^*)v(r^*)_{\perp}$		$\frac{x^3y}{(1-x)^2}$
$F_8 = (rr^*)v(rr^*)_{\perp}'$		$\frac{x^4y}{(1-x)^2}$

Table 3: The eight types of irreducible final partial states and their generating functions.

3.1.5 The f -vector and Euler characteristic of the configuration space $\mathcal{S}_{2,n}$

Theorem 1.4. Let $\mathcal{S}_{2,n}$ be the configuration space for the robot of length n moving in a rectangular tunnel of width 2. If $c_{n,d}$ denotes the number of d -dimensional cubes in $\mathcal{S}_{2,n}$, then

$$C(x, y) = \sum_{n,d \geq 0} c_{n,d} x^n y^d = \frac{1 + xy + x^2 + x^2y + x^3 + 3x^3y + x^3y^2 + 2x^4y + 2x^4y^2 + x^5y^2}{1 - 2x + x^2 - x^3 - x^4 - 2x^4y - 2x^5y - x^5y^2 - x^6y^2}.$$

The above series starts:

$$\begin{aligned} C(x, y) = & 1 + x(y + 2) + x^2(3y + 4) + x^3(y^2 + 8y + 8) + x^4(4y^2 + 18y + 15) \\ & + x^5(11y^2 + 38y + 28) + x^6(y^3 + 30y^2 + 81y + 53) + \dots \end{aligned}$$

in agreement with Table 1.

Proof. This follows from Corollary 3.3 and Propositions 3.4 and 3.5. □

Corollary 3.6. The generating function counting the number c_n of states of $\mathcal{S}_{2,n}$ is

$$\sum_{n \geq 0} c_n x^n = \frac{1 + x^2 + x^3}{1 - 2x + x^2 - x^3 - x^4} = 1 + 2x + 4x^2 + 8x^3 + 15x^4 + 28x^5 + 53x^6 + \dots$$

Proof. This is a straightforward consequence of Theorem 1.4, substituting $y = 0$. □

Theorem 1.5. The Euler characteristic of the configuration space $\mathcal{S}_{2,n}$ equals 1.

Proof. The generating function for the Euler characteristic of $\mathcal{S}_{2,n}$ is

$$\begin{aligned} \sum_{n \geq 0} \chi(\mathcal{S}_{2,n}) x^n &= \sum_{n \geq 0} \left(\sum_{d \geq 0} (-1)^d c_{n,d} \right) x^n = C(x, -1) \\ &= \frac{1 - x - x^3 + x^5}{1 - 2x + x^2 - x^3 + x^4 + x^5 - x^6} \\ &= \frac{1}{1 - x} = 1 + x + x^2 + x^3 + \dots \end{aligned}$$

by Theorem 1.4. All the coefficients of this series are equal to 1, as desired. \square

4 The combinatorics of CAT(0) cube complexes

4.1 CAT(0) spaces

We now define CAT(0) spaces. Consider a metric space X where every two points x and y can be joined by a path of length $d(x, y)$, known as a *geodesic*. Let T be a triangle in X whose sides are geodesics of lengths a, b, c , and let T' be the triangle with the same lengths in the Euclidean plane. For any geodesic chord of length d connecting two points on the boundary of T , there is a comparison chord between the corresponding two points on the boundary of T' , say of length d' . If $d \leq d'$ for any such chord in T , we say that T is a *thin triangle* in X . We say the metric space X has *non-positive global curvature* if every triangle in X is thin.




Figure 8: A chord in a triangle in X , and the corresponding chord in the comparison triangle in \mathbb{R}^2 . The triangle in X is *thin* if $d \leq d'$ for all such chords.

Definition 4.1. A metric space X is said to be CAT(0) if:

- there is a unique geodesic (shortest) path between any two points in X , and
- X has non-positive global curvature.

We are interested in proving that the cubical complex $\mathcal{S}_{m,n}$ is CAT(0), because of the following theorem:

Theorem 4.2. [1, 3, 5, 11] If the configuration space of a robot is a CAT(0) cubical complex, there is an algorithm to find the optimal way of moving the robot from one position to another.

4.2 CAT(0) cubical complexes

It is not clear a priori how one might show that a cubical complex is CAT(0); Definition 4.1 certainly does not provide an efficient way of testing this property. Fortunately, for cubical complexes, we know of two possible approaches.

The topological approach. The first approach uses Gromov’s groundbreaking result that this subtle metric property has a topological–combinatorial characterization:

Theorem 4.3. [9] *A cubical complex is CAT(0) if and only if it is simply connected, and the link of every vertex is a flag simplicial complex.*

Recall that a simplicial complex Δ is *flag* if it has no empty simplices; more explicitly, if the 1-skeleton of a simplex is in Δ , then that simplex must be in Δ . It is easy to see that, in the configuration space of a robot, every vertex has a flag link. Therefore, one approach to prove that these spaces are CAT(0) is to prove they are simply connected; see [1, 8] for examples of this approach.

The combinatorial approach. We will use a purely combinatorial characterization of finite CAT(0) cube complexes [5, 12, 13, 15]; we use the formulation of Ardila, Owen, and Sullivan [5]. After drawing enough CAT(0) cube complexes, one might notice that they look a lot like distributive lattices. In trying to make this statement precise, one discovers a generalization of Birkhoff’s Fundamental Theorem of Distributive Lattices; there are bijections:

$$\begin{aligned} \text{distributive lattices} &\longleftrightarrow \text{posets} \\ \text{rooted CAT(0) cubical complexes} &\longleftrightarrow \text{posets with inconsistent pairs (PIPs)} \end{aligned}$$

Hence to prove that a cubical complex is a CAT(0) cubical complex, it suffices to choose a root for it, and identify the corresponding PIP. By the above correspondence, this PIP will serve as a “certificate” that the complex is CAT(0). For a reasonably small complex, it is straightforward to identify the PIP, as described below. For an infinite family of configuration spaces like the one that interests us, one may do this for a few small examples and hope to identify a pattern that one can prove in general. Along the way, one gains a better understanding of the combinatorial structure of the complexes one is studying. This approach was first carried out in [3] for two examples, and the goal of this paper is to carry it out for the robots $R_{m,n}$. This family is considerably more complicated than the ones in [3], in particular, because it appears to be the first known example where the skeleton of the complex is not a distributive lattice, and the PIP does indeed have inconsistent pairs.




Figure 9: A PIP and the corresponding rooted CAT(0) cubical complex.

Let us describe this method in more detail.

Definition 4.4. [5, 15] A poset with inconsistent pairs (PIP) is a poset P together with a collection of inconsistent pairs, which we denote $p \leftrightarrow q$ (where $p \neq q$), such that

$$\text{if } p \leftrightarrow q \text{ and } q < q' \text{ then } p \leftrightarrow q'.$$

Note that PIPs are equivalent to *event structures* [15] in computer science, and are closely related to *pocsets* [12, 13] in geometric group theory. The *Hasse diagram* of a poset with inconsistent pairs (PIP) is obtained by drawing the poset, and connecting each $<$ -minimal inconsistent pair with a dotted line. The left panel of Figure 9 shows an example.

Theorem 4.5. [5, 12, 13] There is a bijection $P \mapsto X(P)$ between posets with inconsistent pairs and rooted CAT(0) cube complexes.

It is useful to describe both directions of this bijection.

Rooted CAT(0) cubical complex \mapsto PIP: A CAT(0) cube complex X has a system of *hyperplanes* as described by Sageev [13]. Each d -cube C in X has d “hyperplanes” of codimension 1, each one of which is a $(d - 1)$ -cube orthogonal to an edge direction of C and splits it into two equal parts. When two cubes C and C' share an edge e , we identify the two hyperplanes in C and in C' that are orthogonal to e . The result of all these identifications is the set of hyperplanes of X . The right panel of Figure 9 shows a CAT(0) cube complex and its six hyperplanes.

Now we define the PIP P associated to X . The elements of P are the hyperplanes of X . For hyperplanes H_1 and H_2 , we declare $H_1 < H_2$ if, starting at the root v of X , one must cross hyperplane H_1 before one can cross hyperplane H_2 . Finally, we declare $H_1 \leftrightarrow H_2$ if, starting at the root v of X , it is impossible to cross both H_1 and H_2 without backtracking. The left panel of Figure 9 shows the PIP associated to the rooted complex of the right panel.

PIP \mapsto rooted CAT(0) cubical complex: Let P be a PIP. Recall that an *order ideal* of P is a subset I such that if $x < y$ and $y \in I$ then $x \in I$. We say that I is *consistent* if it contains no inconsistent pair.

The vertices of $X(P)$ are identified with the consistent order ideals of P . There is a cube $C(I, M)$ for each pair (I, M) of a consistent order ideal I and a subset $M \subseteq I_{\max}$, where I_{\max} is the set of maximal elements of I . This cube has dimension $|M|$, and its vertices are obtained by removing from I the $2^{|M|}$ possible subsets of M . The cubes are naturally glued along their faces according to their labels. The root is the vertex corresponding to the empty order ideal.

This bijection is also illustrated in Figure 9; the labels of the cubical complex on the right correspond to the consistent order ideals of the PIP on the left.

5 The coral PIP, coral tableaux, and the proof of Theorem 1.2

We have now described all the preliminaries necessary to prove our main result that the configuration space $\mathcal{S}_{m,n}$ of the robotic arm in a tunnel is a CAT(0) cubical complex. We will achieve this by identifying the PIP corresponding to it under the bijection of Theorem 4.5. Interestingly, this PIP has some resemblance with Young’s lattice of partitions. However, instead of partitions, its elements correspond to certain paths which we call *coral snakes*.

5.1 The coral PIP

Definition 5.1. A coral snake λ of height at most m is a path of unit squares, colored alternately black and red (starting with black), inside the tunnel of width m such that:

- (i) The snake λ starts at the bottom left of the tunnel, and takes steps up, down, and right.
- (ii) Suppose λ turns from a vertical segment V_1 to a horizontal segment H to a vertical segment V_2 at corners C_1 and C_2 . Then V_1 and V_2 face the same direction if and only if C_1 and C_2 have the same color. (Note: We consider the first column of the snake a vertical segment going up, even if it consists of a single cell.)

The length $l(\lambda)$ is the number of unit squares of λ , the height $h(\lambda)$ is the number of rows it touches, and the width $w(\lambda)$ is the number of columns it touches. We say that μ contains λ , in which case we write $\lambda \preceq \mu$, if λ is an initial sub-snake of μ obtained by restricting to the first k cells of μ for some k . We write $\lambda \prec \mu$ if $\lambda \preceq \mu$ and $\lambda \neq \mu$.

Figure 10 shows a coral snake λ of length $l(\lambda) = 11$, height $h(\lambda) = 3$, and width $w(\lambda) = 7$. We encourage the reader to check condition (ii). We often omit the colors of the coral snake when we draw them, since they are uniquely determined.




Figure 10: A mathematical and (a photograph of) a real-life coral snake.

Remark 5.2. Our notion of containment of coral snakes differs from the notion of containment in the plane. For example, if λ is the snake with two boxes corresponding to one step right, and μ is the snake with four boxes given by consecutive steps up-right-down, then λ is contained in μ in the plane. However, λ is not a sub-snake of μ and therefore $\lambda \not\preceq \mu$.

Definition 5.3. Define the coral PIP $C_{m,n}$ as follows:

- *Elements:* pairs (λ, s) of a coral snake λ with $h(\lambda) \leq m$ and a non-negative integer s with $s \leq n - l(\lambda) - w(\lambda) + 1$.
- *Order:* $(\lambda, s) \leq (\mu, t)$ if $\lambda \preceq \mu$ and $s \geq t$.
- *Inconsistency:* $(\lambda, s) \leftrightarrow (\mu, t)$ if neither λ nor μ contains the other.

For simplicity, we call the elements of the coral PIP numbered snakes.

Remark 5.4. Figures 4 and 11 illustrate the coral PIPs $C_{m,2}$ for the tunnel of width 2. They have a nice self-similar structure² in the following sense: they grow vertically supported on a main vertical “spine”, forming several sheets of roughly triangular shape. Along the way they grow other vertical spines, and every other spine supports a smaller coral PIP. The situation for higher m is similar, though more complicated.

²During a break, after weeks of unsuccessful attempts to describe these PIPs, the Pacific Ocean sent us a beautiful coral that looked just like them. The proverbial fractal structure of real-life corals helped us discover the self-similar nature of these PIPs; this led to their precise definition and inspired their name.

We will prove the following strengthening of Theorem 1.2.

Theorem 5.5. *The configuration space $\mathcal{S}_{m,n}$ of the robotic arm $R_{m,n}$ of length n in a tunnel of width m is a CAT(0) cubical complex. When it is rooted at the horizontal position of the arm, its corresponding PIP is the coral PIP $C_{m,n}$ of Definition 5.3.*




Figure 11: The PIP $C_{m,n}$ for $m = 2$ and $n = 9$. Some of the labels are omitted; to obtain them, note that each vertical column consists of the elements (λ, s) for a fixed shape λ and $0 \leq s \leq n - l(\lambda) - w(\lambda) + 1$, listed in decreasing order with respect to s .

5.2 Coral tableaux

Although we did not need to mention them explicitly in the statement of Theorem 5.5, certain kinds of tableaux played a crucial role in its discovery, and are indispensable in its proof.

Definition 5.6. *A coral snake tableau (or simply coral tableau) T on a coral snake λ is a filling of the squares of λ with non-negative integers which are strictly increasing horizontally and weakly increasing vertically, following the direction of the snake. We call $\lambda = \text{sh}(T)$ the shape of T , and we say that T is of type (m, n) if $h(\lambda) \leq m$ and $\max(T) + l(\lambda) \leq n$.*

Note that if T is of type (m, n) , it is also of type (m', n') for any $m' \geq m$ and $n' \geq n$.

Definition 5.7. *We call a coral tableau tight if its entries are constant along columns and increase by one along rows.*




Figure 12: Two coral tableaux; the one on the right is tight.

Lemma 5.8. *There is a bijection between tight coral tableaux of type (m, n) and the numbered coral snakes of the PIP $C_{m,n}$.*

Proof. A tight coral tableau T is uniquely determined by its shape and first entry, so the bijection is given by sending T to $(\text{sh}(T), \min(T))$; it suffices to observe that $\max(T) = \min(T) + w(\lambda) - 1$ when T is tight, so the inequality $\min(T) \leq n - l(\lambda) - w(\lambda) + 1$ holds if and only if $\max(T) + l(\lambda) \leq n$. \square

For example, the tight tableau on the right of Figure 12 corresponds to $(\lambda, 2)$ where λ is the snake of Figure 10.

Lemma 5.9. *The possible states of the robotic arm $R_{m,n}$ are in bijection with the coral tableaux of type (m, n) .*

Proof. We encode a position P of the robotic arm as a coral tableau T , building it up from left to right as follows. Every time that the robot takes a vertical step in row i , we add a new square to row i of the tableau, and fill it with the number of the column the step was taken in. An example is shown in Figure 13.

It is clear that the entries of T increase weakly in the vertical direction, and increase strictly in the horizontal direction. It remains to check that the snake $\text{sh}(T)$ is indeed a coral snake. Suppose $\text{sh}(T)$ goes from a vertical V_1 to a horizontal H to a vertical V_2 , turning at corners C_1 and C_2 of row i . Assume for definiteness that V_1 points up and C_1 is black. Then the black and red entries on H represent up and down steps that the robot takes on row i ; so the direction of V_2 is determined by the color of C_2 as desired. \square




Figure 13: From a state R to a coral tableau T .




Figure 14: A tight state and its corresponding tight snake.

We call a state of the robot *tight* if its corresponding snake tableau is tight, see Figure 14.

5.3 Proof of our main structural theorem

We are now ready to prove our strengthening of the main result of this paper, Theorem 1.2.

Theorem 5.5. The configuration space $\mathcal{S}_{m,n}$ of the robotic arm $R_{m,n}$ of length n in a tunnel of width m is a CAT(0) cubical complex. When it is rooted at the horizontal position of the arm, its corresponding PIP is the coral PIP $C_{m,n}$ of Definition 5.3.

Proof. We need to show that, when rooted at the horizontal position, $\mathcal{S}_{m,n}$ is the cubical complex $X(C_{m,n})$ associated to the PIP $C_{m,n}$ under the bijection of Theorem 4.5. We proceed in three steps.

Step 1. Decomposing $\mathcal{S}_{m,n}$ and $C_{m,n}$ by shape.

For a coral snake λ of type (m,n) let $\mathcal{S}_{m,n}^\lambda$ be the induced subcomplex of $\mathcal{S}_{m,n}$ whose vertices are the coral tableaux T with $\text{sh}(T) \preceq \lambda$. Similarly, let $C_{m,n}^\lambda$ be the subPIP of $C_{m,n}$ consisting of the pairs (μ, s) such that $\mu \preceq \lambda$. Notice that $C_{m,n}^\lambda$ is a poset which has no inconsistent pairs. We have

$$\mathcal{S}_{m,n} = \bigcup_{\lambda \text{ of type } (m,n)} \mathcal{S}_{m,n}^\lambda \quad C_{m,n} = \bigcup_{\lambda \text{ of type } (m,n)} C_{m,n}^\lambda \quad (1)$$

where each (non-disjoint) union is over the coral tableaux λ of type (m,n) . (Of course, since $\lambda \preceq \lambda'$ implies $\mathcal{S}_{m,n}^\lambda \subseteq \mathcal{S}_{m,n}^{\lambda'}$ and $C_{m,n}^\lambda \subseteq C_{m,n}^{\lambda'}$, it is sufficient to take the union over those λ which are maximal under inclusion.)

Step 2. Showing $\mathcal{S}_{m,n}^\lambda = X(C_{m,n}^\lambda)$ for each shape λ .

We begin by establishing a bijection between the sets of vertices of both complexes. The vertices of $X(C_{m,n}^\lambda)$ correspond to the consistent order ideals of the coral PIP $C_{m,n}^\lambda$. Since $C_{m,n}^\lambda$ has no inconsistent pairs, these are *all* the order ideals of $C_{m,n}^\lambda$, which are the elements of the distributive lattice $J(C_{m,n}^\lambda)$.

The vertices of $\mathcal{S}_{m,n}^\lambda$ correspond to the coral tableaux T of type (m,n) with shape $\text{sh}(T) \preceq \lambda$ by Lemma 5.9. Let us extend each such tableau T to a tableau of shape λ by adding entries equal to ∞ on each cell of $\lambda - \text{sh}(T)$. We may then identify the set of vertices $V(\mathcal{S}_{m,n}^\lambda)$ with the resulting set of *extended coral tableaux* of shape λ whose finite entries x satisfy $x + l(\lambda) \leq n$.

This set $V(\mathcal{S}_{m,n}^\lambda)$ of extended λ -tableaux forms a poset under reverse componentwise order. Note that if T_1 and T_2 are in $V(\mathcal{S}_{m,n}^\lambda)$ then the componentwise maximum $T_1 \wedge T_2$ and the componentwise minimum $T_1 \vee T_2$ are also in $V(\mathcal{S}_{m,n}^\lambda)$. Then the meet \wedge and join \vee make $V(\mathcal{S}_{m,n}^\lambda)$ into a lattice. In fact, the definitions of \wedge and \vee imply that this lattice is distributive. Birkhoff's Fundamental Theorem of Distributive Lattices then implies that $V(\mathcal{S}_{m,n}^\lambda) \cong J(P)$ where P is the subposet of join-irreducible elements of $V(\mathcal{S}_{m,n}^\lambda)$. One easily verifies that P consists precisely of the tight tableaux of type (m,n) and shape λ , so $P \cong C_{m,n}^\lambda$. It follows that the coral tableaux T of type (m,n) with shape $\text{sh}(T) \preceq \lambda$ are indeed in bijection with the consistent order ideals of the coral PIP $C_{m,n}^\lambda$.

We can make the bijection more explicit. Given a coral tableau T of shape $\mu \preceq \lambda$, we can write $T = T_1 \vee T_2 \vee \cdots \vee T_{l(\mu)}$ as follows. For each i let μ_i be the subsnake consisting of the first i boxes of μ , and let T_i be the unique tight tableau of shape λ_i whose i th entry is equal to the i th entry of T . In fact we can reduce this to a minimal equality

$$T = \bigvee_{i \text{ jump in } T} T_i$$

Remark 5.10. One might hope that Theorem 5.5 could be generalized for a robotic arm moving in a d -dimensional tunnel $[0, n_1] \times \cdots \times [0, n_d]$. However, the resulting cubical complexes are not CAT(0) in general. Even in the simplest 3-dimensional case $[0, n] \times [0, 1] \times [0, 1]$ the result does not generalize. Figure 17 illustrates two examples of the cubical complexes for $n = 1$ and $n = 2$. The case $n = 1$ consists of three vertices (states) and three edges (allowable moves) forming a triangle (without the interior face), while the case $n = 2$ consists of seven vertices and eight edges forming an hexagon and a triangle glued together along an edge. In both cases, these cubical complexes are not CAT(0). For instance, they have non-contractible loops and there is not a unique geodesic connecting two opposite vertices of the hexagon.




Figure 17: The cubical complex of a robotic arm in a 3-dimensional tunnel $[0, n] \times [0, 1] \times [0, 1]$ is not CAT(0) for $n = 1$ or $n = 2$.

6 Shortest path and the diameter of the transition graph

The coral tableau representation turns out to be very useful for finding the distance between any two possible states of the robot in the transition graph $G(R_{m,n})$, as well as for finding its diameter. Before carrying this program out in detail, let us briefly describe the intuition behind it.

By Theorem 5.5, a position of the robot corresponds to a consistent order ideal in the coral PIP $C_{m,n}$ of Figures 4 and 11. As we mentioned in Remark 5.4 and these figures illustrate, these PIPs are obtained by gluing several *sheets* of roughly triangular shape along some vertical spines; in the proof of Theorem 5.5, these sheets are the subposets $C_{m,n}^\lambda$ for the maximal λ . Every such sheet grows out of the main spine, possibly branching along the way. A careful look at the inconsistent pairs shows that any consistent order ideal in this PIP must be contained within a single sheet.

Now suppose we wish to find the shortest path between two states P and P' of the robot. This is equivalent to finding the shortest way of transforming an order ideal I in a sheet T of the PIP into an order ideal I' in another sheet T' . To do this, one must first shrink the ideals I and I' into ideals $J \subseteq I$ and $J' \subseteq I'$ which lie in the intersection $T \cap T'$ of the sheets, and then find the best way of transforming J into J' inside $T \cap T'$. The following definitions make this more precise.

F 
$$\begin{aligned} v &= (21, 20, 19, 17, 16) \\ w &= (14, 12, 10, 8, 7, 5) \end{aligned}$$
$$\begin{aligned} v' &= (22, 19, 17, 14, 11) \\ w' &= (10, 7, 5, 3, 1) \end{aligned}$$
$$d(P, P') = |w| + |v - v'| + |w'|, \quad (3)$$

We will prove this proposition in Section 6.1. We will also use it to find an explicit formula for the diameter of the transition graph $G(R_{m,n})$.

Lemma 6.3. *For $m = 1$, the maximum distance between two positions of the robot in the transition graph $G(R_{1,n})$ is attained by the pair $(urdrurdr \dots, rr \dots r)$ of left justified and fully horizontal robotic arms. The diameter of $G(R_{1,n})$ is*

For larger m the situation is rather different. Define the *left justified robot* $L_{m,n}$, the *shifted left justified robot* $L_{m,n}^+$, and the *fully horizontal robot* $H_{m,n}$ of length n as

respectively. The first two of these are illustrated in Figure 19.




Figure 19: The left justified robot $L_{3,23}$ and the shifted left justified robot $L_{3,23}$.

Theorem 6.4. *The diameter of the transition graph $G(R_{m,n})$ of a robot on length n in a tunnel of width $m \geq 2$ is*

$$\text{diam } G(R_{m,n}) = \begin{cases} d(L_{m,n}, H_{m,n}) & \text{for } n < 6 \\ d(L_{m,n}, L_{m,n}^+) & \text{for } n \geq 6. \end{cases}$$

These distances are explicitly given by

$$\begin{aligned} d(L_{m,n}, H_{m,n}) &= s_n - z_{m,n} \\ d(L_{m,n}, L_{m,n}^+) &= s_{n-1} - z_{m,n} + s_{n-2} - z_{m,m+n-3}, \end{aligned}$$

where $s_n = \binom{n+1}{2}$ and $z_{m,n} = (r+1)k + \binom{k}{2}(m+1)$ with $n = (m+1)k + r$, $0 \leq r \leq m$.

This answers a question of the first author at the Open Problem Session of the Encuentro Colombiano de Combinatoria ECCO 2014 [2].

6.1 Proof of Proposition 6.2

Let T and T' be two coral tableaux of fixed type (m, n) , and let λ and λ' their corresponding shapes. We decompose T into two parts T_1 and T_2 of shapes $\lambda \wedge \lambda'$ and $\lambda \setminus (\lambda \wedge \lambda')$ respectively. Similarly, we decompose T' into two parts T'_1 and T'_2 of shapes $\lambda \wedge \lambda'$ and $\lambda' \setminus (\lambda \wedge \lambda')$ respectively. Additionally, we create another filling of λ starting with the number n in the bottom left square and decreasing each time by one along the snake. The restriction of this filling to the shape $\lambda \setminus (\lambda \wedge \lambda')$ is denoted by \overline{T}_2 . The filling \overline{T}_2' of $\lambda' \setminus (\lambda \wedge \lambda')$ is defined analogously. These constructions are illustrated in Figure 20.




Figure 20: Decomposition of two coral tableaux T and T' of type $(3, 23)$.

Given two fillings T and T' of the same shape λ , we denote by $T - T'$ the filling of λ whose entries are the differences between the entries in T and T' . Note that some of the entries of $T - T'$ might be negative. Recall that $|T - T'|$ is the sum of the absolute values of $T - T'$.

Lemma 6.5. *Let P and P' be two positions of the robotic arm $R_{m,n}$, and T and T' be the corresponding coral tableaux. Using the notation above, the distance between P and P' in the transition graph $G(R_{m,n})$ is equal to*

$$d(P, P') = |\overline{T_2} - T_2| + |T_1 - T'_1| + |\overline{T'_2} - T'_2|.$$

In the example of Figure 20, $d(P, P') = 56 + 12 + 26 = 94$. We will show a shortest path subdivided into three parts connecting $T \rightarrow T_1 \rightarrow T'_1 \rightarrow T'$ of lengths 56, 12, and 26.

Proof. First, we prove that the number of moves needed to move the robot P to P' is at least the claimed number. For this we analyze the possible moves of the robot in terms of the corresponding coral tableau. There are two kinds of moves: switching a corner of the robot corresponds to either increasing or decreasing an entry of the coral tableau by one, such that the result is still a coral tableau. Flipping the end of the robot corresponds to either deleting or adding a last box of the tableau when its entry has the maximum possible value $n - l(\lambda)$, in agreement with Definition 5.6. For simplicity, we call these steps *allowable tableau moves*.

Let λ and λ' be the shapes of T and T' respectively. Let $T_1, T_2, \overline{T_2}$ and $T'_1, T'_2, \overline{T'_2}$ as above. In order to move from T to T' with allowable tableau moves it is necessary to:

- (i) make disappear all entries of T_2 in T ,
- (ii) convert all the entries of T_1 to the entries of T'_1 , and
- (iii) make appear all entries in T'_2 . We claim that the number of moves required to do these three steps is at least $|\overline{T_2} - T_2|$, $|T_1 - T'_1|$ and $|\overline{T'_2} - T'_2|$ respectively, from which we deduce that the distance between T and T' is at least

$$|\overline{T_2} - T_2| + |T_1 - T'_1| + |\overline{T'_2} - T'_2|.$$

To prove the first claim, let t and $\bar{t} = n - l(\lambda) + 1$ be the last entries of T_2 and $\overline{T_2}$, respectively. To make the last cell of T_2 disappear, we need to perform $n - l(\lambda) - t$ allowable tableau moves at that entry to increase it to its maximum value of $n - l(\lambda)$, and one additional move to remove the cell, for a total of $|\bar{t} - t|$ moves. Continuing analogously, we see that if we wish to achieve (i) we need to perform at least $|\overline{T_2} - T_2|$ tableau moves in $\lambda' - (\lambda \wedge \lambda')$. Similarly, to achieve (iii) we need at least $|\overline{T'_2} - T'_2|$ in $\lambda - (\lambda \wedge \lambda')$. Finally, since each allowable move changes an entry by one, we require at least $|T_1 - T'_1|$ moves in $(\lambda \wedge \lambda')$ to carry out (ii).

The previous argument is also a roadmap for how to achieve this lower bound. We first go from tableau T to T_1 in $|\overline{T_2} - T_2|$ moves by removing the boxes of T_2 in order from the last to the first. Then we then go from T_1 to $\max\{T_1, T'_1\}$ by increasing the entries one at a time from the last to the first box, and from $\max\{T_1, T'_1\}$ to T'_1 by decreasing the entries one at a time, again from the last to the first box; this takes a total of $|T_1 - T'_1|$ moves. Finally, we connect T'_1 to T' by adding all the entries in T'_2 from first to last, using $|\overline{T'_2} - T'_2|$ moves. \square

Proposition 6.2 is now a direct consequence of the previous lemma.

Proof of Proposition 6.2. Let P and P' be two positions of the robotic arm $R_{m,n}$, and T and T' be their corresponding coral tableaux. We need to show that

$$|w| = |\overline{T_2} - T_2|, \quad |v - v'| = |T_1 - T'_1|, \quad |w'| = |\overline{T'_2} - T'_2|.$$

The entry of a box in $\overline{T_2} - T_2$ counts the number of links in P after the corresponding vertical link, including it. Therefore, the entries of w are exactly equal to the entries of $\overline{T_2} - T_2$. Similarly, the entries of w' are equal to the entries of $\overline{T'_2} - T'_2$. Now, if $v = (v_1, \dots, v_k)$ and if (t_1, \dots, t_k) are the entries of T_1 in the order they appear along the snake, then $v_i = n - (t_i + i - 1)$ for each i . Therefore the entries of $v - v'$ are exactly the negatives of the entries of $T_1 - T'_1$, and hence $|v - v'| = |T_1 - T'_1|$. \square

6.2 Proof of Theorem 6.4

Throughout this section we fix $m \geq 2$ and $n \in \mathbb{N}$. We say that a coral snake λ is of type (m, n) if it appears in the coral PIP $C_{m,n}$; that is, if $h(\lambda) \leq m$ and $l(\lambda) + w(\lambda) - 1 \leq n$. All the coral snakes and positions of the robot considered in this section are of type (m, n) .

Definition 6.6. *The shape of a position P of the robot, denoted $\text{shape}(P)$, is the shape of its corresponding coral tableau. The intersection shape of two positions P and P' is defined as $\text{shape}(P) \wedge \text{shape}(P')$.*

Our strategy to find the diameter of the transition graph $G(R_{m,n})$ is to find the maximum distance between two positions of the robot with a fixed intersection shape λ , and then maximize this quantity over all shapes λ . We need to distinguish two kinds of snakes.

Definition 6.7. *A coral snake of type (m, n) is said to be a:*

- *side snake: if the end point of any robot with coral tableau of shape λ is on one of the two horizontal sides of the tunnel.*
- *middle snake: if it is not a side snake.*

Definition 6.8. *We define the following special positions of the robot (see Figure 21):*

- L_λ : *the position with componentwise maximum vertical labelling among all positions whose coral tableau contains the shape λ .*
- L_λ^+ : *the position with componentwise maximum vertical labelling among all positions whose coral tableau contains the shape λ and whose intersection shape with L_λ is λ .*
- H_λ : *the position with componentwise minimum vertical labelling among all robots whose coral tableau has shape λ .*




Figure 21: Four special positions of the robot for the given λ , $(m, n) = (3, 23)$ and $k = 16$.

The following two technical lemmas are the key steps to proving Theorem 6.4.

Lemma 6.9. *The maximum distance between two positions of the robot with fixed intersection shape λ is:*

- (i) $d(L_\lambda, H_\lambda)$, if λ is a side snake.
- (ii) $\max\{d(L_\lambda, H_\lambda), d(L_\lambda, L_\lambda^+)\}$, if λ is a middle snake.

Lemma 6.10. *The following hold:*

- (i) If $\lambda \prec \lambda'$, then $d(L_\lambda, H_\lambda) > d(L_{\lambda'}, H_{\lambda'})$.
- (ii) If $\lambda \prec \lambda'$ are two middle snakes, then $d(L_\lambda, L_\lambda^+) > d(L_{\lambda'}, L_{\lambda'}^+)$.

We postpone the proof of these two lemmas for the moment and use them to prove our main diameter result.

Proof of Theorem 6.4. By Lemma 6.9, the diameter of the transition graph $G(R_{m,n})$ is the maximum between the two values

$$\max_{\lambda} d(L_\lambda, H_\lambda), \quad \max_{\text{middle snakes } \lambda} d(L_\lambda, L_\lambda^+).$$

By Lemma 6.10, we have

$$\max_{\lambda} d(L_\lambda, H_\lambda) = d(L_\emptyset, H_\emptyset), \quad \max_{\text{middle snakes } \lambda} d(L_\lambda, L_\lambda^+) = d(L_\square, L_\square^+).$$

where \emptyset is the empty snake and \square is the snake that consists of only one box. By definition, we have that $L_\emptyset = L_\square = L_{m,n}$, $H_\emptyset = H_{m,n}$ and $L_\square^+ = L_{m,n}^+$. Therefore,

$$\text{diam } G(R_{m,n}) = \max\{d(L_{m,n}, H_{m,n}), d(L_{m,n}, L_{m,n}^+)\}.$$

The explicit formulas for these distances stated in Theorem 6.4 are obtained directly from Proposition 6.2. For instance, $d(L_{m,n}, H_{m,n})$ is the sum of all vertical labels in $L_{m,n}$. This sum is equal to the sum $s_n = 1 + \dots + n$ of all labels (including the horizontal ones) minus the sum $z_{m,n}$ of the horizontal labels. The formula for $d(L_{m,n}, L_{m,n}^+)$ is obtained similarly, considering the labels of $L_{m,n}^+$ as well.

When $n \geq 4$, we may rewrite $d(L_{m,n}, L_{m,n}^+) = d(L_{m,n}, H_{m,n}) + s_{n-4} - z_{m,n-4} - 2$, which shows that $d(L_{m,n}, L_{m,n}^+) \geq d(L_{m,n}, H_{m,n})$ for $n \geq 6$ and $d(L_{m,n}, L_{m,n}^+) < d(L_{m,n}, H_{m,n})$ for $n = 4, 5$. The cases $n = 1, 2, 3$ are easily verified by hand. \square

6.2.1 Fixing a shape: proof of Lemma 6.9

Let us prove Lemma 6.9 for any fixed coral snake λ of type (m, n) . Let P and P' be two positions of the robot with intersection shape λ . We begin by proving case (ii), which is more intricate, and then return to case (i).

Proof of Lemma 6.9 (ii). Let λ be a middle snake. For $\ell(\lambda) + w(\lambda) - 1 \leq k \leq n$ the following special positions of the robot (see Figure 21) will play an important role in the proof:

- $L_{\lambda,k}$: the position with componentwise maximum vertical labelling among all positions whose entries are less than or equal to k and whose coral tableau contains the shape λ .

- $L_{\lambda,k}^+$: the position with componentwise maximum vertical labelling among all positions whose entries are less than or equal to k , whose coral tableau contains the shape λ , and whose intersection shape with $L_{\lambda,k}$ is λ .

Notice that $L_{\lambda,k}$ (resp. $L_{\lambda,k}^+$) consists of $n - k$ horizontal steps followed by the position analogous to L_λ (resp. L_λ^+) for the robotic arm of length k . In particular, $L_\lambda = L_{\lambda,n}$ and $L_\lambda^+ = L_{\lambda,n}^+$. The description of the maximum distance between two robots with fixed intersection shape λ will follow from the following three lemmas.

Lemma 6.11. *For any positions P and P' with intersection shape λ there exists k such that $d(P, P') \leq \max\{d(L_\lambda, L_{\lambda,k}^+), d(L_\lambda^+, L_{\lambda,k})\}$.*

Proof. Recall that

$$d(P, P') = |w| + |v - v'| + |w'|.$$

We will perform two distance-increasing transformations to the pair (P, P') to obtain either $(L_\lambda, L_{\lambda,k}^+)$ or $(L_\lambda^+, L_{\lambda,k})$. Without loss of generality assume that, among the vertical labels in v and v' , the minimum m is achieved in v' .

The first transformation moves the vertical steps corresponding to v in P as far to the left as possible (by increasing their values in v), and moves the vertical steps corresponding to v' in P' as far to the right as possible (by decreasing their values in v') while keeping the minimal label equal to m . Denote by P_1 and P'_1 the resulting positions; Figure 22 shows an example. This transformation preserves the values of $|w|$ and $|w'|$ in the formula and increases the value of $|v - v'|$.

The second transformation changes the parts of P_1 and P'_1 after their intersection λ , in order to maximize $|w|$ and $|w'|$ while keeping $|v - v'|$ fixed. We need to keep the intersection shape equal to λ , so we must respect the up/down direction of the next vertical step in P_1 and P'_1 after λ ; we convert the rest of P_1 and P'_1 into “zigzag snakes” subject to that constraint. The result is either the pair $(L_\lambda, L_{\lambda,k}^+)$ or $(L_\lambda^+, L_{\lambda,k})$ for some k . It follows that

$$d(P, P') \leq d(P_1, P'_1) \leq \max\left\{d(L_\lambda, L_{\lambda,k}^+), d(L_\lambda^+, L_{\lambda,k})\right\},$$

as desired. This construction is illustrated in Figure 22. \square

Lemma 6.12. $d(L_\lambda, L_{\lambda,k}^+) \geq d(L_\lambda^+, L_{\lambda,k})$.

Proof. For simplicity, let $\ell = n - k$ be the number of initial horizontal steps of $L_{\lambda,k}$ and $L_{\lambda,k}^+$ (if λ is non-empty). As above, label the links of the robot in decreasing order from n up to 1 along the shape of the robot. Let h, h^+, h_k, h_k^+ be the restriction of this labelling to the horizontal links located after the last vertical link corresponding to λ in each of the robots $L_\lambda, L_\lambda^+, L_{\lambda,k}, L_{\lambda,k}^+$ respectively.

The vector $h - h_k$ (resp. $h^+ - h_k^+$) consists of an initial sequence of $(n - k)$ s followed by the entries of h (resp. h^+) that are less than $n - k$. Consider the injection taking the i th horizontal step of L_λ to the i th horizontal step of $L_{\lambda,k}^+$, which has a weakly larger label. Under this injection, each term of $h - h_k$ is dominated by the corresponding term in $h^+ - h_k^+$. It follows that $(|h^+ - h_k^+|) - (|h - h_k|) \geq 0$. \square

Lemma 6.13. *For a fixed middle snake λ , $d(L_\lambda, L_{\lambda,k}^+)$ is concave up as a function of k , and therefore attains its maximum value either at $k = l(\lambda) + w(\lambda) - 1$ or at $k = n$.*




Figure 22: Increasing the distance between two robots with a given intersection shape.

Proof. It suffices to show that the (possibly negative) difference $d(L_\lambda, L_{\lambda, k+1}^+) - d(L_\lambda, L_{\lambda, k}^+)$ is increasing for every k where it is defined. Let $v_k^+ \circ w_k^+$ be the $(L_\lambda, L_{\lambda, k}^+)$ -decomposition of $L_{\lambda, k}^+$, and $\ell(w_k^+)$ be the number of entries in w_k^+ . If $\ell(\lambda)$ is the number of boxes of the intersection shape, then

$$d(L_\lambda, L_{\lambda, k+1}^+) - d(L_\lambda, L_{\lambda, k}^+) = \ell(\lambda) - \ell(w_k^+).$$

The result follows from the fact that $\ell(w_k^+)$ is decreasing in k and $\ell(\lambda)$ is constant. \square

We finally have all the ingredients to complete the proof of Lemma 6.9(ii). Let λ be a middle snake. By Lemmas 6.11 and 6.12,

$$d(P, P') \leq d(L_\lambda, L_{\lambda, k}^+)$$

for some value k . By Lemma 6.13, the maximum of $d(L_\lambda, L_{\lambda, k}^+)$ is achieved either at $k = n$ or $k = \ell(\lambda) + w(\lambda) - 1$. These two cases correspond precisely to $d(L_\lambda, L_\lambda^+)$ and $d(L_\lambda, H_\lambda)$ respectively. Therefore,

$$d(P, P') \leq \max\{d(L_\lambda, H_\lambda), d(L_\lambda, L_\lambda^+)\}$$

as we wished to show. \square

Now we go back to case (i), which is easier.

Proof of Lemma 6.9 (i). If λ is a side snake, then either $\text{shape}(P) = \lambda$ or $\text{shape}(P') = \lambda$, otherwise their intersection shape would be bigger than λ . For simplicity assume $\text{shape}(P') = \lambda$. Again, we transform P and P' to increase their distance: we move the vertical steps corresponding to λ in P as far to the left as possible, and the vertical steps of P' as far to the right as possible. This transformation preserves the values of $|w|$ and $|w'|$ in the formula and increases the value of $|v - v'|$. In addition, transform the part of the modified P after λ in order to maximize $|w|$. The result is the pair (L_λ, H_λ) . \square

6.2.2 Varying the shape: proof of Lemma 6.10

Proof of Lemma 6.10 (i). By Proposition 6.2, removing the last box of λ increases the distance $d(L_\lambda, H_\lambda)$ by $\ell(\lambda)$. Therefore $\lambda \prec \lambda'$ implies $d(L_\lambda, H_\lambda) > d(L_{\lambda'}, H_{\lambda'})$. \square

To prove part (ii) we use a basic lemma.

Lemma 6.14. *Let A be a subsequence of $(p, \dots, 2, 1)$ obtained by removing an arithmetic progression with common difference c , and let A' be a subsequence of $(p', \dots, 2, 1)$ obtained by removing an arithmetic progression with common difference c . If $p > p'$ then $|A| \geq |A'|$.*

Proof. Say A and A' have a and a' elements, respectively. When listed in decreasing order, the largest a' terms of A are greater than or equal to the a' terms of A' . The remaining $a - a'$ terms (if there are any) are positive. Therefore $|A| \geq |A'|$. \square

Proof of Lemma 6.10 (ii). Assume that $\lambda \prec \lambda'$ are two middle snakes. We need to show that $d(L_\lambda, L_{\lambda'}^+) = |w| + |v - v'| + |w'|$ decreases as we increase λ . The term $|v - v'| = 0$ stays constant. The term $|w|$ clearly does not decrease when we make λ larger. Finally, the term $|w'|$ also does not decrease when we make λ thanks to Lemma 6.14. The desired result follows. \square

7 Implementation of the shortest path algorithm

Since the configuration space of the robotic arm $R_{m,n}$ is $\text{CAT}(0)$, we are able to navigate it thanks to the following theorem.

Theorem 4.2. [1, 3, 11] If the configuration space of a robot is a $\text{CAT}(0)$ cubical complex, there is an explicit algorithm to move the robot optimally from one position to another, in:

- *the edge metric:* the number of moves, if only one move at a time is allowed,
- *the cube metric:* the number of steps (where in each step we may perform several physically independent moves),
- *the time metric:* the time elapsed if we may perform independent moves simultaneously.

As shown in [5], there is also a (more complicated) algorithm to move the robot optimally in terms of the Euclidean metric in the configuration space. However, this metric is less relevant, as it does not seem to have a natural interpretation in terms of the robot.

The algorithms of Theorem 4.2 are described in detail in [3]; they are based on the *normal cube paths* of Niblo and Reeves [10]. We have implemented these algorithms in Python for the robotic arms discussed in this paper. To use the program, the user inputs the width of the tunnel and two positions of the robotic arm of the same length, expressed as sequences of steps of the form: r (right), u (up), d (down) steps. The program outputs the distance between the two states in the edge metric and in the cube metric³, and an animation taking the robot between the two states. The downloadable code, instructions, and a sample animation are at <http://math.sfsu.edu/federico/Articles/movingrobots.html>.

³It turns out that the optimal paths in the cube metric are also optimal in the time metric.

8 Future directions and open problems

- Our robotic arms $R_{m,n}$ can never have links facing left. Now consider the robotic arm of length n in a tunnel of width m whose links may face up, down, right, **or left**, starting in the fully horizontal position, and subject to the same kinds of moves: switching corners and flipping the end. Is the configuration space of this robot also CAT(0)? In [4], we show the answer is affirmative for width $m = 2$. We believe this is also true for larger m , but proving it in general seems to require new ideas.
- Theorem 6.4 gives the diameter of the configuration space $\mathcal{S}_{m,n}$ in the edge metric, that is, the largest number of moves separating two positions of the robot $R_{m,n}$. What is the diameter of $\mathcal{S}_{m,n}$ in the cube metric, that is, the largest number of steps separating two positions of the robot $R_{m,n}$, where a step may consist of several physically independent moves?

9 Acknowledgements

This paper includes results from HB’s undergraduate thesis at Universidad del Valle in Cali, Colombia (advised by FA and CC) and JG’s undergraduate project at San Francisco State University in San Francisco, California (advised by FA). We thank the SFSU-Colombia Combinatorics Initiative, which made this collaboration possible. We would also like to thank the Pacific Ocean for bringing us the beautiful coral that inspired Definition 5.3 and the proof of our main Theorem 1.6.

References

- [1] A. Abrams and R. Ghrist. State complexes for metamorphic robots. *Int. J. Robotics Res.* **23** (2004) 811-826.
- [2] F. Ardila. Open Problem Session, Encuentro Colombiano de Combinatoria, Bogotá, 2014.
- [3] F. Ardila, T. Baker, and R. Yatchak. Moving robots efficiently using the combinatorics of CAT(0) cubical complexes. *Adv. in Appl. Math.* **48** (2012) 142-163.
- [4] F. Ardila, H. Bastidas, C. Ceballos, and J. Guo. The configuration space of a robotic arm in a tunnel of width 2. To appear in the Proceedings of Formal Power Series and Algebraic Combinatorics (FPSAC 2016), *Discrete Mathematics and Theoretical Computer Science*.
- [5] F. Ardila, M. Owen, and S. Sullivant. Geodesics in CAT(0) cubical complexes. *SIAM J. Discrete Math.* **28-2** (2014) 986-1007.
- [6] G. Birkhoff. Rings of sets. *Duke Mathematical Journal* **3** (1937) 443-454.
- [7] M. Bridson and A. Haefliger. Metric spaces of non-positive curvature. Springer-Verlag, Berlin, Heidelberg, 1999.
- [8] R. Ghrist and V. Peterson. The geometry and topology of reconfiguration. *Adv. in Appl. Math.* **38** (2007) 302-323.

- [9] M. Gromov. Hyperbolic groups. In Essays in group theory, volume 8 of *Math. Sci. Res. Inst. Publ.* 75-263. Springer, New York, 1987.
- [10] G.A. Niblo and L.D. Reeves. The geometry of cube complexes and the complexity of their fundamental groups. *Topology*, **37**(3) (1998) 621-633.
- [11] L. D. Reeves. Biautomatic structures and combinatorics for cube complexes. *Ph.D. thesis, University of Melbourne*, 1995.
- [12] M. A. Roller. Poc sets, median algebras and group actions. an extended study of Dunwoody's construction and Sageev's theorem. Unpublished preprint, 1998.
- [13] M. Sageev. Ends of group pairs and non-positively curved cube complexes. *Proc. London Math. Soc.* (3), **71** (3) 585-617, 1995.
- [14] L. Santocanale. A nice labelling for tree-like event structures of degree 3. *Inf. Comput.*, 208:652-665, 2010.
- [15] G. Winskel. Event Structures. *Advances in Petri Nets 1986*. Springer Lecture Notes in Computer Science 255, 1987.